



ATTORNEY DOCKET NO. 5437-055

BOX NEW APPLICATION
THE ASSISTANT COMMISSIONER FOR PATENTS
WASHINGTON, D.C. 20231

SIR:

TRANSMITTED HERewith FOR FILING IS THE PATENT APPLICATION OF
INVENTOR(S) CHARLES MERRIAM

FOR: **METHOD AND APPARATUS FOR MANAGING INFORMATION RETENTION SYSTEM**

ENCLOSED ARE:

- ☒ 18 PAGES SPECIFICATION, 33 CLAIMS (9 Pages), 1 SHEET ABSTRACT, 6 SHEETS OF DRAWINGS, AND 1 COVER PAGE.
- ☒ DECLARATION AND POWER OF ATTORNEY (Executed) (3 pgs)
- ☒ AN ASSIGNMENT OF THE INVENTION TO Sun Microsystems, Inc.

THE FILING FEE IS CALCULATED BELOW

CLAIMS	(1) FOR	(2) NUMBER FILED	(3) NUMBER EXTRA	(4) RATE	(5) CALCULATIONS
TOTAL CLAIMS	33	-20 =	13	X \$ 18 =	\$ 234.00
INDEPENDENT	3	- 3 =	0	X \$ 78 =	\$.00
MULTIPLE DEPENDENT CLAIMS (if applicable)				+ \$260 =	\$ -0-
LATE FILING OF DECLARATION				+ \$130 =	\$ -0-
BASIC FEE					\$ 690.00
TOTAL OF ABOVE CALCULATIONS =					\$ 924.00
Reduction by 50% for filing by Small Entity					\$.00
<input type="checkbox"/> FILING IN NON-ENGLISH LANGUAGE				+ \$130 =	\$.00
<input checked="" type="checkbox"/> RECORDATION OF ASSIGNMENT				+ \$ 40 =	\$ 40.00
TOTAL					\$ 964.00

- ☒ A CHECK IN THE AMOUNT OF \$964.00 TO COVER THE FILING FEE IS ENCLOSED.
- ☒ THE ASSISTANT COMMISSIONER IS HEREBY AUTHORIZED TO CHARGE ANY ADDITIONAL FEES WHICH MAY BE REQUIRED FOR THE PAPERS BEING FILED HERewith AND FOR WHICH NO CHECK IS ENCLOSED HERewith, OR CREDIT ANY OVERPAYMENT TO DEPOSIT ACCOUNT NO. 50-0448. A DUPLICATE COPY OF THIS SHEET IS ENCLOSED.

SABATH & TRUONG

111 No. Market Street
Suite 815
San Jose, California 95113
Tel: (408) 293-9934
Fax: (408) 293-2183

BOBBY K. TRUONG
Attorney of Record
Reg. No. 37,499

"Express Mail" mailing number EL433967628US

Date of Deposit February 9, 2000

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Box New Application, Assistant Commissioner for Patents, Washington, D.C. 20231.

Annette Jacobs
Name

Signature

February 9, 2000
Date

METHOD AND SYSTEM FOR MANAGING INFORMATION RETENTION

INVENTOR: Charles Merriam
a citizen of United State of America
residing at
1447 Flicker Way
Sunnyvale, California 94087

Assignee: Sun Microsystems, Inc.
901 San Antonio Road
MS PALO1-521
Palo Alto, California 94303

Entity Status: Large

METHOD AND SYSTEM FOR MANAGING INFORMATION RETENTION

Inventor(s): Charles Merriam

Background

5 This invention relates generally to information systems, and more particularly to a method and apparatus for managing information retention.

 To limit the amount of information that is maintained in an information system, many companies implement policies (often referred to as information retention policies) under which information sets are systematically deleted from the
10 system. These policies set forth criteria for determining if and when an information set is to be deleted, and are often time-based, such that an information set is deleted after having been in the system for a certain amount of time, or condition-based, such that an information set is deleted after a certain condition has been satisfied, or a combination of both. For example, a set of information may be deleted if it has been
15 in the system for a certain amount of time AND it has not been accessed within a certain period of time. By systematically deleting information sets from the system, the system is maintained at a reasonable size. This helps to limit the amount of storage required to implement the system, and the effort needed to maintain the system.

20 Typically, an information retention policy is enforced by deleting from the system information sets which have been determined under the retention policy as being eligible for deletion. An information set may be deleted by removing references to the information set, thereby rendering the information set inaccessible, or by overwriting the information set with new information, or both. While removing an
25 information set from a system by way of deletion is sound in principle, it is difficult to

implement in practice. One of the main problems with the deletion method is that in order to delete an instance of an information set, it is first necessary to find it. If there are multiple copies of an information set (as is often the case if the system is large or distributed or if backup copies are made), it is necessary to find and to delete all
5 copies of the information set before the information set is completely removed from the information system. For a large information system, it can be a tedious, daunting, and sometimes impossible task to find all copies of an information set. As a result, stray copies of an information set often linger in the system even after the information set has supposedly been deleted. These stray copies can seriously undermine the
10 effectiveness of the retention policy. As this discussion illustrates, the current methodology for enforcing information retention policies leaves much to be desired. As a result, an improved mechanism for managing information retention is needed.

Summary of the Invention

15 The present invention provides an improved mechanism for managing information retention in an information system. The present invention is based, as least partially, upon the observation that it is not necessary to delete an information set from a system in order to purge it. Instead, all that is necessary is to make the information set unrenderable to a user. If the information set is made unrenderable,
20 then it is for all intents and purposes useless; as a result, it is effectively "purged" from the system, even if it remains physically in the system.

In light of this observation, the present invention provides an information retention management mechanism, whereby, whenever an information set enters an information system, a key is associated with the information set. The information set
25 is encrypted using the associated key, and the encrypted form of the information set is

stored in the information system. The unencrypted form of the information set is not stored. Then, at a later time, the information set is rendered to a user by accessing the encrypted form of the information set, accessing the associated key, and decrypting the encrypted information set using the associated key to derive the original information set. Once derived, the information set is rendered to the user. So long as the associated key remains in the system, this process may be carried out to render the information set to a user.

At some point, in accordance with an information retention policy, the information set is selected for purging. To purge the information set, all that needs to be done is to delete the associated key. Once that is done, all copies of the encrypted information set stored within the information system are made unrenderable; as a result, the information set is effectively "purged" from the system. Notice that it is not necessary to delete the encrypted information set from the system. The information set is made unrenderable, and hence, purged even if the encrypted information set remains physically in the information system. Thus, the present invention eliminates the need to find and to delete each and every copy of the information set. By doing so, the present invention transforms the potentially daunting and highly difficult task of finding and deleting every instance of an information set into a simple one of deleting a key. Not only does this simplify the information retention process, but it also increases the effectiveness of the process because it solves the "stray copy" problem. As noted above, even if a copy of the encrypted information set remains in the system, no harm is done because the encrypted information set cannot be rendered. Consequently, the present invention provides a significantly improved information retention mechanism.

Brief Description of the Drawings

Fig. 1 is a block diagram of a system in which the present invention may be implemented.

Fig. 2 is a flow diagram showing the operation of the information manager in response to an information set from an information source.

Fig. 3 is a flow diagram showing the operation of the retention manager.

Fig. 4 is a flow diagram showing one embodiment of the operation of the information manager in providing an information set to an information sink.

Fig. 5 is a flow diagram showing an alternative embodiment of the operation of the information manager in providing an information set to an information sink.

Fig. 6 is a block diagram of a hardware computer system in which various components of the present invention may be implemented.

Detailed Description of the Embodiment(s)

With reference to Fig. 1, there is shown a block diagram of a system in which the present invention may be implemented, the system comprising an information system 102, an information source 104, and an information sink 106. For the sake of simplicity, only one information source 104 and one information sink 106 are shown; however, it should be noted that multiple information sources and information sinks may interact with the information system 102. For purposes of the present invention, the information source 104 represents any mechanism capable of originating and providing one or more information sets to the information system 102. Similarly, the information sink 106 represents any mechanism capable of receiving one or more information sets from the information system 102, and rendering the information set to a user. In this regard, information source 104 and information sink 106 may be, but

are not limited to being, general purpose computers running client software. The information source 104 and information sink 106 may be coupled to the information system 102 via a direct connection or via a network.

It should be noted that the information source 104 and information sink 106 may be the same mechanism. At one point, a mechanism may provide an information set to the information system 102, and hence, act as an information source. At another time, that same mechanism may receive an information set from the information system 102, and hence, act as an information sink. Thus, the same mechanism may be capable of performing both roles.

The information system 102 is the overall entity responsible for storing and maintaining information sets. As used herein, the term "information set" refers broadly to any type of information, including but not limited to files, messages, web pages, communications, cryptographic keys, access codes, etc. This information may take on many different forms, including but not limited to text, graphics, audio/video, data, etc. All such forms are within the scope of the present invention. Essentially, any type of information that can be stored and retrieved may be managed by information system 102.

The information system 102 may take on many different forms. For example, information system 102 may be a file server responsible for managing access to files stored within one or more repositories. In such an implementation, the information source 104 and information sink 106 would be clients of the file server. The information system 102 may also be a web server responsible for providing web pages to clients in response to access requests. In such a case, the information source 104 would be a content providing mechanism (e.g. an ftp client), and the information sink 106 would be a web browser. In addition, the information system 102 may also be a

mail server responsible for managing access to electronic mail (email) messages. In such an implementation, the information source 104 and information sink 106 would be mail clients of the mail server. These and other implementations of the information system 102 are within the scope of the present invention.

5 Whatever form the information system 102 takes, it comprises an information manager 110, one or more information repositories 112, a key repository 114, and a retention manager 116. The information manager 110 is primarily responsible for interfacing with the information source 104 and the information sink 106. More specifically, the information manager 110 receives one or more information sets from
10 the information source 104, and stores them into one or more of the information repositories 112. In addition, the information manager 110 accesses one or more information sets from the repositories 112, and provides them to the information sink 106. In performing these tasks, the information manager 110 processes the information sets in accordance with the methodology of the present invention. In one
15 embodiment, the information manager 110 comprises an encryption engine 118. As will be elaborated upon in a later section, the encryption engine 118 is used to encrypt information sets prior to storing them into the repositories 112, and to decrypt information sets prior to providing them to the information sink 106. For purposes of the invention, the encryption engine 118 may implement any known encryption
20 methodology, including but not limited to single key encryption and symmetric key encryption. In one embodiment, the information manager 110 is implemented as a set of instructions executable by one or more processors. However, if so desired, the encryption engine 118 and, optionally, the information manager 110 as a whole may be implemented using special purpose, hard-wired logic components. This and other
25 implementations are within the scope of the present invention.

006020 "REF0560

The keys used by the information manager 110 to encrypt and to decrypt the information sets are maintained in the key repository 114. Depending upon the encryption methodology implemented by the encryption engine 118, the form of the keys may differ. For example, the keys may be single keys or symmetric key pairs (e.g. public/private keys). In one embodiment, each key or key set has a key ID associated therewith. These ID's enables the keys to be easily referenced, accessed, and manipulated.

The content of the key repository 114 is managed by the retention manager 116. More specifically, the retention manager 116 is responsible for implementing an information retention policy to determine which keys to delete from the key repository 114. As will be explained in a later section, deleting a key from the key repository 114 in effect purges from the information system 102 all information sets associated with that key. Thus, by managing the keys in the key repository 114, the retention manager 116 in effect manages the retention of information sets within the information system 102. In performing this management function, the retention manager 116 may implement any desired retention policy. For example, retention manager 116 may implement a time-based policy whereby keys are deleted from the key repository 114 after having been in the system 102 for a certain amount of time. The retention manager 116 may also implement a condition-based policy, whereby keys are deleted upon the satisfaction of one or more conditions. In addition, retention manager 116 may implement a classification-based policy, whereby all keys having a certain classification are deleted at a certain time. These and other policies (or combinations thereof) may be implemented by the retention manager 116 within the spirit of the invention. In one embodiment, the retention manager 116 is implemented as a set of instructions executable by one or more processors. However,

if so desired, the retention manager 116 may be implemented using special purpose, hard-wired logic components. This and other implementations are within the scope of the present invention.

5 With reference to the system diagram of Fig. 1 and the flow diagrams of Figs. 2-4, the operation of the information system 102 will now be described. Fig. 2 shows the operation of the information manager 110 when it receives and stores an information set from an information source 104. Fig. 3 shows the operation of the retention manager 116 in managing the content of the key repository 116, and Fig. 4
10 illustrates the operation of the information manager 110 when it accesses and provides an information set to an information sink 104.

 Referring first to Fig. 2, whenever the information manager 110 receives (202) an information set from an information source 104, it makes a determination (204) as to whether a key has already been associated with that information set. In one
15 embodiment, this is done by checking the header of the information set for a key ID. If no key ID is found, then it means that the information set has not been previously stored in the information system 102. In such a case, the information manager 110 selects (212) a key from the key repository 114 to associate with this information set. This key selection may be performed based upon any predetermined policy. For
20 example, the key may be selected based upon the current date so that the key of the day is selected, or the selection may be based upon the classification of the information set (e.g. highly confidential, confidential, and generally available). For purposes of the present invention, any key selection policy may be implemented.

 Once the key (having a key ID associated therewith) is selected, the
25 information manager 110 encrypts (214) the information set using the selected key,

and then stores (216) the encrypted information set 120 into one or more of the persistent information repositories 112. As part of the storage process, the information manager 110 writes the key ID into the header (which remains unencrypted) of the encrypted information set 120. This key ID associates the selected key with the encrypted information set 120, and enables the proper key to be accessed at a later time from the key repository 114 to allow the encrypted information set 120 to be decrypted. An information set is thus persistently stored within the information system 102.

An important point to note is that only the encrypted form of the information set 120 is persistently stored within the information system 102. The unencrypted information set does not get persistently stored. This is important in preserving the integrity of the information retention policy implemented by the retention manager 116, as will be discussed below. It is understood that in order to encrypt the information set, the information manager 110 needs to have, at some point, at least a portion of the unencrypted information set in volatile, non-persistent memory. This does not pose a problem. So long as the unencrypted information set is not stored in persistent storage, the information retention policy will not be undermined.

Returning now to (204) of Fig. 2, if it is determined that a key has already been associated with the information set (as may be the case if the information set had already been stored in the information system 102, and had been provided to the information source 104), then the information manager 110 extracts the key ID from the header of the information set, and uses the key ID to access (206) the associated key in the key repository 114. Once the associated key is accessed, the information manager 110 encrypts (208) the information set using the associated key, and stores (210) the encrypted information set 120 into one or more of the persistent information

repositories 112. As part of the storage process, the information manager 110 re-
writes the key ID into the header of the encrypted information set 120. The encrypted
information set 120 is thus stored within the information system 102. As noted
previously, only the encrypted form 120 of the information set is persistently stored
5 within the information system 102. The unencrypted information set is not
persistently stored.

Once the encrypted information set 120 is stored within the information
system 102, it can be managed by the retention manager 116. More specifically, the
retention manager 116 can control whether the encrypted information set 120 is
10 retained or purged from the system 102. For purposes of the present invention, the
term "purge" is defined broadly to mean making a set of information unrenderable to a
user. Under this definition, it is not necessary to delete an information set from a
system in order to purge it. Rather, all that is necessary is to make the information set
unrenderable to a user. If the information set is made unrenderable, then it is for all
15 intents and purposes useless, so that it is effectively "purged" from the system even if
it remains physically within the system. In the present invention, information sets are
stored in the information system 102 only in encrypted form. In such a system, it is
an easy matter to make an information set unrenderable. All that needs to be done is
to delete the key needed to decrypt the encrypted information set. Once that is done,
20 all instances of the encrypted information set are made unrenderable. Thus, the
present invention transforms the task of managing information retention into one of
managing keys. By managing the keys in the key repository 114, the retention
manager 116 controls whether information sets are retained or purged from the
information system 102.

Referring now to Fig. 3, the retention manager 116 operates by first implementing (302) a predetermined information retention policy. This policy sets forth the criteria for determining if and when an information set is to be purged from the information system 102. For example, this retention policy may be time-based, so that information sets are purged after having been in the system 102 for a certain amount of time, or condition-based, so that information sets are purge upon the satisfaction of one or more conditions, or classification-based, so that information sets having a certain classification are purged at a certain time. These and other retention policies (or combinations thereof) may be implemented by the retention manager 116.

Based upon the retention policy, the retention manager 116 determines (304) whether any information sets need to be purged from the information system 102. If none need to be purged, then the retention manager 116 loops back to (302). However, if one or more information sets need to be purged, then the retention manager 116 deletes (306) from the key repository 114 the key(s) associated with the encrypted versions of those information set(s) 120. By doing so, the retention manager 116 makes those encrypted information sets 120 unrenderable, thereby, effectively purging them from the system 102. Notice that it is not necessary to delete the encrypted information sets 120 (although it may be desirable to do so at some point to free up storage space). By eliminating the need to delete information sets from the system, the present invention eliminates the need to find each and every instance of an information set. This serves to greatly simplify the information retention management process.

In addition to storing information sets into the system 102, the information manager 110 also manages access to information sets already stored within the system 102. In one embodiment, the information manager 110 carries out this management

task in accordance with the flow diagram of Fig. 4. Initially, the information manager 110 receives (402) a request from an information sink 106 to access a specific information set within the information system 102. In response to this request, the information manager 110 accesses (404) from the one or more repositories 112 the encrypted form 120 of the requested information set, and then determines (406) the key associated with the encrypted information set 120. In one embodiment, this is done by extracting a key ID from the header of the encrypted information set 120. Armed with the key ID, the information manager 110 searches the key repository 114 to determine (408) whether the key associated with the key ID is available. If the key repository 114 does not have a key associated with the key ID, then it means that the associated key has been deleted. This in turn means that the requested information set was previously purged by the retention manager 116. If that is the case, then the information manager 110 returns (410) an error message to the information sink 106 to inform the information sink 106 that the requested information set has been purged from the system 102 and hence, can no longer be rendered. In this manner, the purging of the requested information set is manifested.

Returning to (408), if the information manager 110 determines that a key associated with the key ID is available, then it accesses (412) the associated key in the key repository 114. Armed with the key, the information manager 110 decrypts (414) the encrypted information set 120 to derive the original information set. The information manager 110 then provides (416) the unencrypted information set to the information sink 106 for rendering to a user. An information set is thus accessed from the information system 102 and rendered. In accordance with one embodiment: (1) the encrypted information set 120 is decrypted only when it is necessary to render the information set to a user; and (2) the information sink 106 comprises sufficient logic

to prevent it from persistently storing the unencrypted information set. These considerations help to prevent the undermining of the retention policy.

Thus far, the information manager 110 has been described as being the component responsible for decrypting the encrypted information set 120. As an alternative, the decryption function may be delegated to the information sink 106.

Fig. 5 is a flow diagram showing the operation of the information manager 110 where the decryption task has been delegated to the information sink 106. As shown, the information manager 110 initially receives (502) a request from the information sink 106 for a specific information set within the information system 102. In response to this request, the information manager 110 accesses the encrypted form 120 of the requested information set from the one or more repositories 112, and provides (504) the encrypted information set 120 to the information sink 106. Included with the encrypted information set 120 is the unencrypted header which contains a key ID. Since it is the encrypted form of the information set that is provided, the information sink 106 may freely store the encrypted information set in its own persistent storage.

At some point, the information sink 106 may need to render the information set to a user. To do so, the information sink 106 needs the key associated with the encrypted information set. The information sink 106 acquires this key by: (1) extracting the key ID from the header of the encrypted information set; and (2) sending a key request, which includes the extracted key ID, to the information manager 110. Upon receiving (506) this request, the information manager 110 searches the key repository 114 to determine (508) whether the key associated with the key ID is available. If the key repository 114 does not have a key associated with the key ID, then it means that the associated key has been deleted. This in turn means that the encrypted information set was previously purged by the retention manager

116. If that is the case, then the information manager 110 returns (510) an error message to the information sink 106 to inform the information sink 106 that the encrypted information set has been purged from the system 102 and hence, can no longer be rendered.

5 On the other hand, if the information manager 110 determines (508) that a key associated with the key ID is available, then it accesses (512) the associated key in the key repository 114, and provides (514) it to the information sink 106. Armed with the key, the information sink 106 decrypts the encrypted information set and renders it to the user. In accordance with one embodiment: (1) the encrypted information set 120
10 is decrypted only when it is necessary to render the information set to the user; and (2) the information sink 106 comprises sufficient logic to prevent it from persistently storing the key received from the information manager 110. These considerations help to prevent the undermining of the retention policy. In the manner described, the present invention provides an improved mechanism for managing information
15 retention in an information system.

Hardware Overview

Fig. 6 is a block diagram that illustrates a computer system 600 in which an embodiment of the invention may be implemented. Computer system 600 includes a
20 bus 602 or other communication mechanism for communicating information, and a processor 604 coupled with bus 602 for processing information. Computer system 600 also includes a main memory 606, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 602 for storing information and instructions to be executed by processor 604. In addition, main memory 606 may be
25 further used for storing temporary variables or other intermediate information during

006020 " REF ID: A560

execution of instructions by processor 604. Computer system 600 further includes a read only memory (ROM) 608 or other static storage device coupled to bus 602 for storing static information and instructions for processor 604. A storage device 610, such as a magnetic disk or optical disk, is provided and coupled to bus 602 for storing
5 information and instructions.

Computer system 600 may be coupled via bus 602 to a display 612, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 614, including alphanumeric and other keys, is coupled to bus 602 for communicating information and command selections to processor 604. Another type
10 of user input device is cursor control 616, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 604 and for controlling cursor movement on display 612. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

15 According to one embodiment, the functionality of the information manager 110 and the retention manager 116 of the present invention are provided by computer system 600 in response to processor 604 executing one or more sequences of one or more instructions contained in main memory 606. Such instructions may be read into main memory 606 from another computer-readable medium, such as storage device
20 610. Execution of the sequences of instructions contained in main memory 606 causes processor 604 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and
25 software.

006020 "SEE FIG 50"

The term "computer-readable medium" as used herein refers to any medium that participates in providing instructions to processor 604 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 610. Volatile media includes dynamic memory, such as main memory 606. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 602. Transmission media can also take the form of acoustic or electromagnetic waves, such as those generated during radio-wave, infra-red, and optical data communications.

Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 604 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 600 can receive the data on the telephone line and use an infra-red transmitter to convert the data to an infra-red signal. An infra-red detector can receive the data carried in the infra-red signal and appropriate circuitry can place the data on bus 602. Bus 602 carries the data to main memory 606, from which processor 604

retrieves and executes the instructions. The instructions received by main memory 606 may optionally be stored on storage device 610 either before or after execution by processor 604.

Computer system 600 also includes a communication interface 618 coupled to bus 602. Communication interface 618 provides a two-way data communication coupling to a network link 620 that is connected to a local network 622. For example, communication interface 618 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 618 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 618 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

Network link 620 typically provides data communication through one or more networks to other data devices. For example, network link 620 may provide a connection through local network 622 to a host computer 624 or to data equipment operated by an Internet Service Provider (ISP) 626. ISP 626 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 628. Local network 622 and Internet 628 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 620 and through communication interface 618, which carry the digital data to and from computer system 600, are exemplary forms of carrier waves transporting the information.

Computer system 600 can send messages and receive data, including program code, through the network(s), network link 620 and communication interface 618. In the Internet example, a server 630 might transmit a requested code for an application program through Internet 628, ISP 626, local network 622 and communication interface 5 618. The received code may be executed by processor 604 as it is received, and/or stored in storage device 610, or other non-volatile storage for later execution. In this manner, computer system 600 may obtain application code in the form of a carrier wave.

10 At this point, it should be noted that although the invention has been described with reference to specific embodiments, it should not be construed to be so limited. Various modifications can be made by those of ordinary skill in the art with the benefit of this disclosure without departing from the spirit of the invention. Thus, the invention should not be limited by the specific embodiments used to illustrate it but 15 only by the scope of the appended claims.

What is claimed is:

1. A method for managing information retention in a system, comprising:
receiving a set of information into a system;
associating one or more keys with said set of information;
5 encrypting said set of information using said one or more keys;
storing said set of information in encrypted form into one or more repositories;
and
purging said set of information from the system by deleting said one or more
keys, thereby making said set of information unrenderable.

10

2. The method of claim 1, wherein said set of information is purged from
the system without requiring that the encrypted form of said set of information be
deleted from the one or more repositories.

15

3. The method of claim 1, wherein said set of information is stored in the
one or more repositories only in encrypted form.

20

4. The method of claim 1, wherein said one or more keys comprises a
symmetrically paired set of keys.

5. The method of claim 1, further comprising:

prior to deletion of said one or more keys, receiving a request from an
information sink to render said set of information to a user;

accessing the encrypted form of said set of information from the one or more
25 repositories;

decrypting the encrypted form of said set of information using said one or more keys to derive said set of information; and

providing said set of information to the information sink to enable the information sink to render said set of information to the user.

5

6. The method of claim 5, wherein said set of information is stored in the one or more repositories only in encrypted form, and wherein the encrypted form of said set of information is decrypted only when it is necessary to render said set of information to the user.

10

7. The method of claim 1, further comprising:
prior to deletion of said one or more keys, receiving a request from an information sink to render said set of information to a user;
accessing the encrypted form of said set of information from the one or more repositories;

15

accessing said one or more keys; and
providing the encrypted form of said set of information and said one or more keys to the information sink to enable the information sink to decrypt the encrypted form of said set of information using said one or more keys to render said set of information to the user.

20

8. The method of claim 7, wherein said set of information is stored in the one or more repositories only in encrypted form, and wherein the encrypted form of said set of information is decrypted only when it is necessary to render said set of information to the user.

25

006020" REF050

9. The method of claim 1, wherein purging comprises:
determining, based upon an information retention policy, whether said set of
information should be purged from the system; and

5 in response to a determination that said set of information should be purged
from the system, purging said set of information from the system by deleting said one
or more keys, thereby making said set of information unrenderable.

10. The method of claim 9, wherein said information retention policy is
10 time-based such that said set of information is purged after a certain period of time.

11. The method of claim 9, wherein said information retention policy is
condition-based such that said set of information is purged when one or more
conditions are satisfied.

15 12. An apparatus for managing information retention in a system,
comprising:

a mechanism for receiving a set of information into a system;

a mechanism for associating one or more keys with said set of information;

20 a mechanism for encrypting said set of information using said one or more
keys;

a mechanism for storing said set of information in encrypted form into one or
more repositories; and

a mechanism for purging said set of information from the system by deleting
25 said one or more keys, thereby making said set of information unrenderable.

13. The apparatus of claim 12, wherein said set of information is purged from the system without requiring that the encrypted form of said set of information be deleted from the one or more repositories.

5

14. The apparatus of claim 12, wherein said set of information is stored in the one or more repositories only in encrypted form.

15. The apparatus of claim 12, wherein said one or more keys comprises a
10 symmetrically paired set of keys.

16. The apparatus of claim 12, further comprising:
a mechanism for receiving, prior to deletion of said one or more keys, a
request from an information sink to render said set of information to a user;
15 a mechanism for accessing the encrypted form of said set of information from
the one or more repositories;
a mechanism for decrypting the encrypted form of said set of information
using said one or more keys to derive said set of information; and
a mechanism for providing said set of information to the information sink to
20 enable the information sink to render said set of information to the user.

17. The apparatus of claim 16, wherein said set of information is stored in
the one or more repositories only in encrypted form, and wherein the encrypted form
of said set of information is decrypted only when it is necessary to render said set of
25 information to the user.

006020 "SECRET" 050

18. The apparatus of claim 12, further comprising:

a mechanism for receiving, prior to deletion of said one or more keys, a request from an information sink to render said set of information to a user;

5 a mechanism for accessing the encrypted form of said set of information from the one or more repositories;

a mechanism for accessing said one or more keys; and

a mechanism for providing the encrypted form of said set of information and said one or more keys to the information sink to enable the information sink to
10 decrypt the encrypted form of said set of information using said one or more keys to render said set of information to the user.

19. The apparatus of claim 18, wherein said set of information is stored in the one or more repositories only in encrypted form, and wherein the encrypted form
15 of said set of information is decrypted by the information sink only when it is necessary to render said set of information to the user.

20. The apparatus of claim 12, wherein the mechanism for purging comprises:

20 a mechanism for determining, based upon an information retention policy, whether said set of information should be purged from the system; and

a mechanism for deleting, in response to a determination that said set of information should be purged from the system, said one or more keys, thereby making said set of information unrenderable.

25

21. The apparatus of claim 20, wherein said information retention policy is time-based such that said set of information is purged after a certain period of time.

22. The apparatus of claim 20, wherein said information retention policy is
5 condition-based such that said set of information is purged when one or more conditions are satisfied.

23. A computer readable medium having stored thereon instructions which, when executed by one or more processors, cause the one or more processors to
10 manage information retention in a system, comprising:

instructions for causing one or more processors to receive a set of information into a system;

instructions for causing one or more processors to associate one or more keys with said set of information;

15 instructions for causing one or more processors to encrypt said set of information using said one or more keys;

instructions for causing one or more processors to store said set of information in encrypted form into one or more repositories; and

instructions for causing one or more processors to purge said set of
20 information from the system by deleting said one or more keys, thereby making said set of information unrenderable.

24. The computer readable medium of claim 23, wherein said set of information is purged from the system without requiring that the encrypted form of
25 said set of information be deleted from the one or more repositories.

25. The computer readable medium of claim 23, wherein said set of information is stored in the one or more repositories only in encrypted form.

26. The computer readable medium of claim 23, wherein said one or more keys comprises a symmetrically paired set of keys.

27. The computer readable medium of claim 23, further comprising:
instructions for causing one or more processors to receive, prior to deletion of
said one or more keys, a request from an information sink to render said set of
information to a user;

instructions for causing one or more processors to access the encrypted form
of said set of information from the one or more repositories;

instructions for causing one or more processors to decrypt the encrypted form
of said set of information using said one or more keys to derive said set of
information; and

instructions for causing one or more processors to provide said set of
information to the information sink to enable the information sink to render said set of
information to the user.

28. The computer readable medium of claim 27, wherein said set of
information is stored in the one or more repositories only in encrypted form, and
wherein the encrypted form of said set of information is decrypted only when it is
necessary to render said set of information to the user.

006020 "SECRET" 050

29. The computer readable medium of claim 23, further comprising:
instructions for causing one or more processors to receive, prior to deletion of
said one or more keys, a request from an information sink to render said set of
information to a user;

5 instructions for causing one or more processors to access the encrypted form
of said set of information from the one or more repositories;

instructions for causing one or more processors to access said one or more
keys; and

instructions for causing one or more processors to provide the encrypted form
10 of said set of information and said one or more keys to the information sink to enable
the information sink to decrypt the encrypted form of said set of information using
said one or more keys to render said set of information to the user.

30. The computer readable medium of claim 29, wherein said set of
15 information is stored in the one or more repositories only in encrypted form, and
wherein the encrypted form of said set of information is decrypted by the information
sink only when it is necessary to render said set of information to the user.

31. The computer readable medium of claim 23, wherein the instructions
20 for causing one or more processors to purge said set of information from the system
comprises:

instructions for causing one or more processors to determine, based upon an
information retention policy, whether said set of information should be purged from
the system; and

instructions for causing one or more processors to delete, in response to a determination that said set of information should be purged from the system, said one or more keys, thereby making said set of information unrenderable.

5 32. The computer readable medium of claim 31, wherein said information retention policy is time-based such that said set of information is purged after a certain period of time.


10 33. The computer readable medium of claim 31, wherein said information retention policy is condition-based such that said set of information is purged when one or more conditions are satisfied.

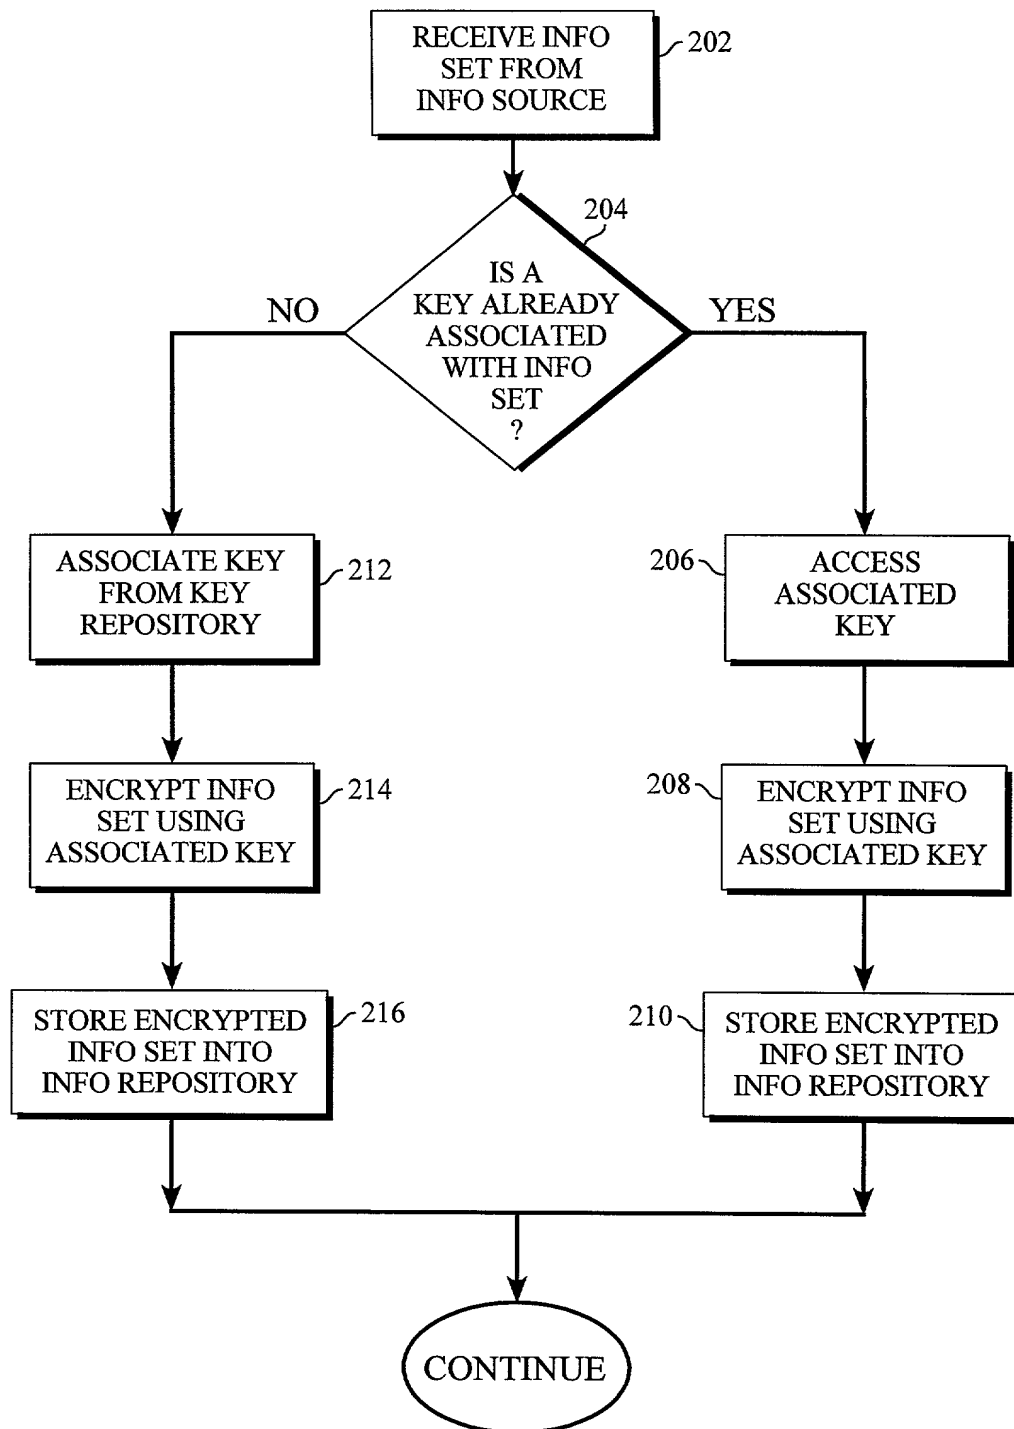
005020 "SECRET" 050

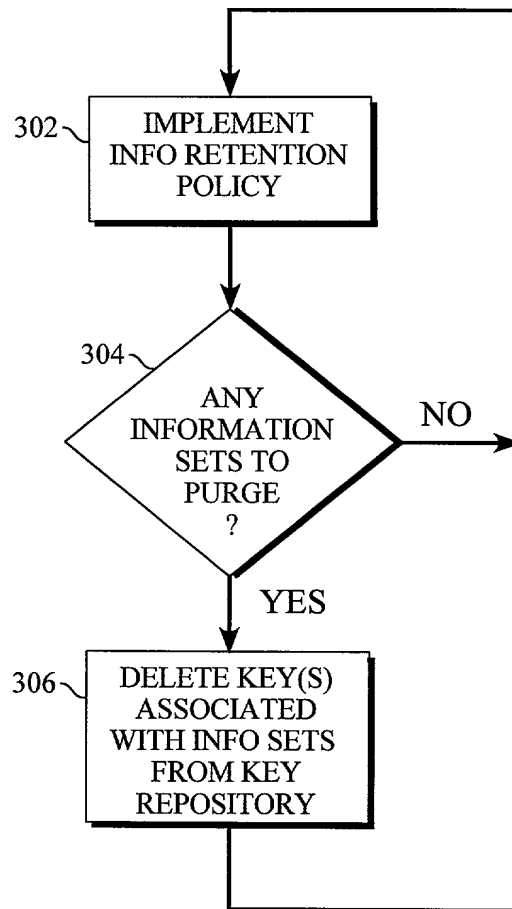
ABSTRACT OF THE DISCLOSURE

An improved information retention management mechanism is disclosed wherein an information set may be purged from an information system without having to delete the information set from the system. Whenever an information set enters an information system, a key is associated with the information set. The information set is encrypted using the associated key, and the encrypted form of the information set is stored in the information system. The unencrypted form of the information set is not stored. To render the information set to a user, the encrypted form of the information set is accessed along with the associated key, and then decrypted using the associated key to derive the original information set. Once derived, the information set is rendered to the user. So long as the associated key remains in the system, this process may be carried out to render the information set to a user. At some point, in accordance with an information retention policy, the information set is selected for purging. To purge the information set, all that needs to be done is to delete the associated key. By deleting the associated key, all copies of the encrypted information set stored within the information system are made unrenderable; as a result, the information set is effectively "purged" from the system. This purging is achieved without having to delete the encrypted information set from the system.

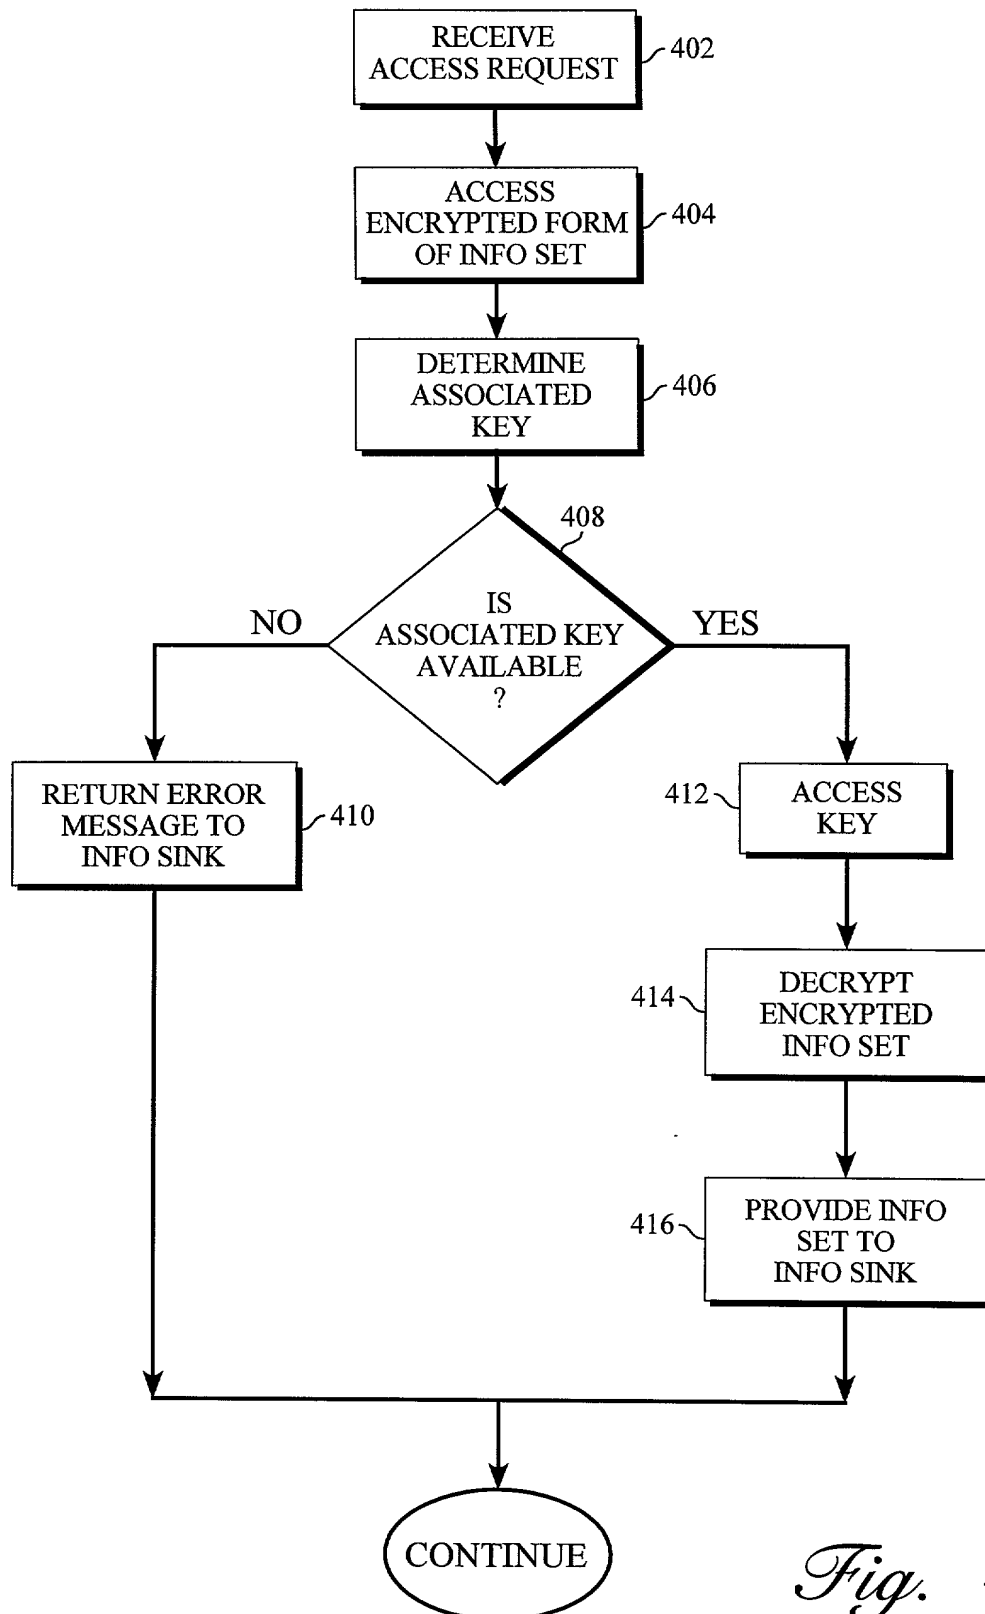
Fig. 1

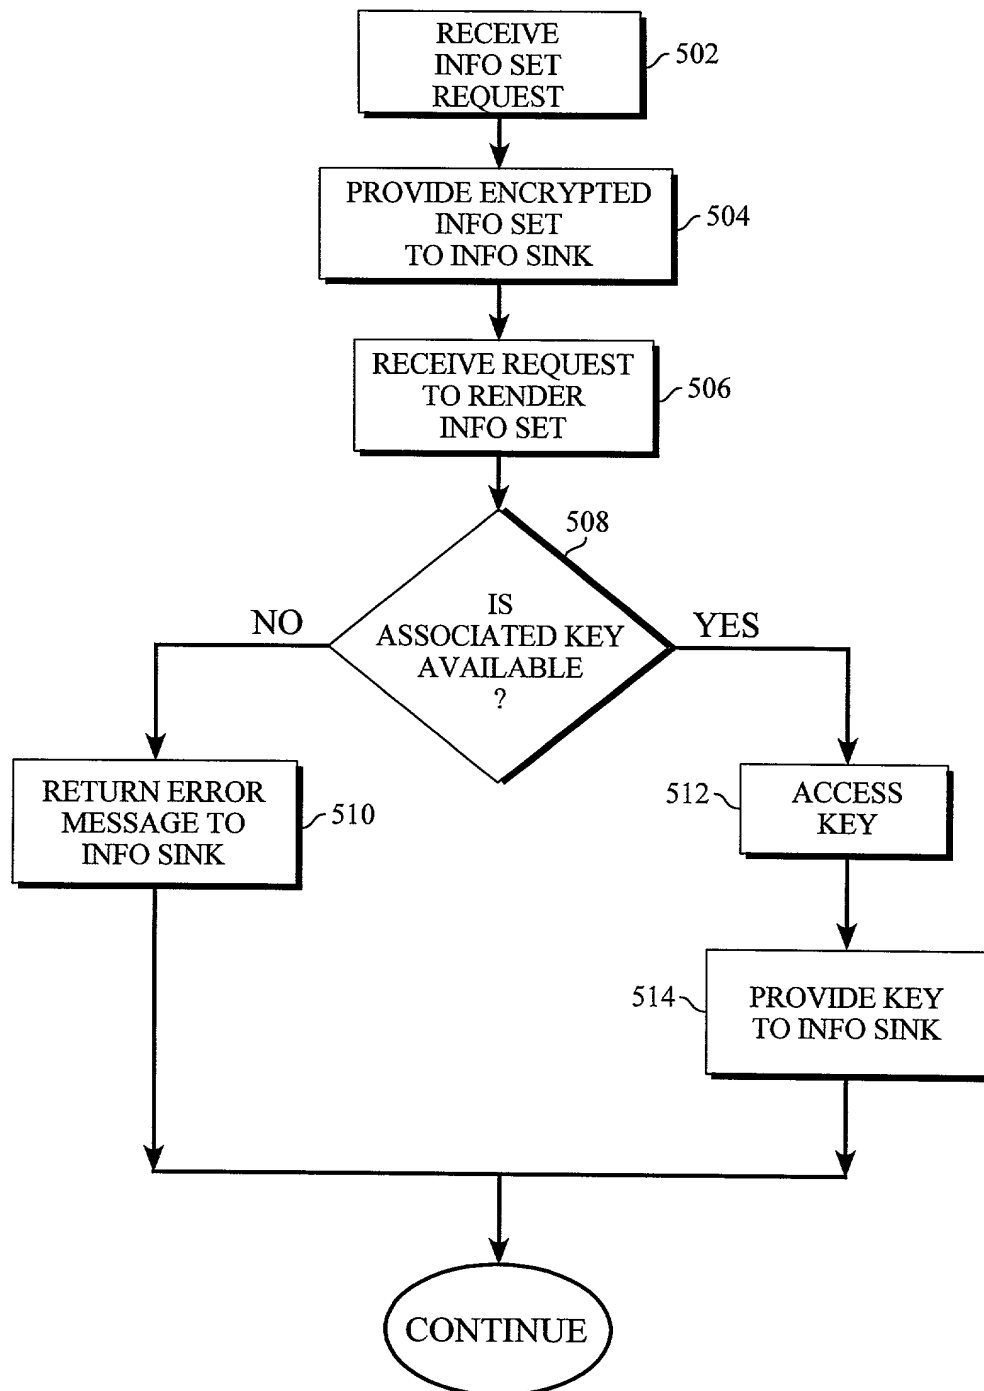


*Fig. 2*

*Fig. 3*

4/6

*Fig. 4*

*Fig. 5*




Fig. 6

PATENT APPLICATION

DECLARATION AND POWER OF ATTORNEY
Original Application

As below named inventor, I declare that I have reviewed and understand the contents of the specification, including the claims, as amended by any amendment specifically referred to in this Declaration, that the information given herein is true, that I believe that I am the original and first inventor of the invention entitled:

METHOD AND SYSTEM FOR MANAGING INFORMATION RETENTION

which is described and claimed in:

- X the attached specification or
- the specification in application Serial No. _____ filed _____.
- The present application is a continuation-in-part of Prior Application Serial No. _____ filed _____, and may be considered to disclose and claim subject matter in addition to that disclosed in the Prior Application, and I hereby claim the benefit of 35 U.S.C. Section 120.

that I acknowledge my duty to disclose information in accordance with 37 C.F.R. Section 1.56 and defined on the attached sheet, which is material to the examination of this application, that I do not know and do not believe the same was ever known or used in the United States of America before my or our invention thereof or patented or described in any printed publication in any country before my or our invention thereof, or more than one year prior to this application, or in public use or on sale in the United States of America more than one year prior to this application, that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representatives or assigns more than twelve months prior to this application and that as to applications for patent or inventor's certificate filed by me or my legal representatives or assigns in any country foreign to the United States of America, the earliest filed foreign application(s) filed within twelve months prior to the filing date of this application and all foreign applications filed more than twelve months prior to the filing date of this application, if any, are identified below.

CHECK APPROPRIATE BOX

- X No earlier-filed foreign applications.
- Required information as to foreign applications filed prior to filing date of this application is on page 5 attached hereto and made a part hereof.

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

<u>NAME</u>	<u>REGISTRATION NO.</u>	<u>NAME</u>	<u>REGISTRATION NO.</u>
Kenneth Olsen	26,493	Bobby K. Truong	37,499
Timothy J. Crean	37,116	John F. Schipper	26,994
Robert S. Hauser	37,847	Stanley N. Protigal	28,657
Joseph T. FitzGerald	33,881	Richard E. Bee	18,005
Alexander E. Silverman	37,940	George M. Steres	36,690
Christine S. Lam	37,489	of SABATH & TRUONG	
Anirma Rakshpal Gupta,	38,275		
Sean Patrick Lewis	42,798		
Michael J. Schallop	44,319		
Bernice B. Chen	42,403		
Kenta Suzue	45,145		
Noreen A. Krall	39,734		
Richard J. Lutton, Jr.	39,756		
of SUN MICROSYSTEMS, INC.			

SEND CORRESPONDENCE TO


Bobby K. Truong, Esq.
SABATH & TRUONG
111 N. Market Street, Suite 815
San Jose, California 95113

DIRECT TELEPHONE CALLS TO

Bobby K. Truong, Esq.
408/293-9934
Fax: 408/293-2183

(201) FULL NAME OF INVENTOR	LAST NAME MERRIAM	FIRST NAME CHARLES	MIDDLE NAME	
RESIDENCE & CITIZENSHIP	CITY SUNNYVALE	STATE OR FOREIGN COUNTRY CALIFORNIA	COUNTRY OF CITIZENSHIP USA	
POST OFFICE ADDRESS	POST OFFICE ADDRESS 1447 FLICKER WAY	CITY SUNNYVALE	STATE OR COUNTRY CA	ZIP CODE 94087

I further declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

Name (201) CHARLES MERRIAM	Signature 	Date 2/8/2000
--------------------------------------	--	-------------------------

Section 1.56 Duty to Disclose Information Material to Patentability.

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclose information exists with respect to each pending claim until the claim is cancelled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is cancelled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by Sections 1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applications to carefully examine:

- (1) prior art cited in search reports of a foreign patent office in a counterpart application, and
- (2) the closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made of record in the application, and

- (1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
- (2) It refutes, or is inconsistent with, a position the application takes in:
 - (i) opposing an argument of unpatentability relied on by the Office, or
 - (ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any considerations given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

- (1) Each inventor named in the application;
- (2) Each attorney or agent who prepares or prosecutes the application; and
- (3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent or inventor.